

IN THE CLAIMS:

Please amend the claims as follows:

1 1. (Currently Amended) In a file server having a storage operating system, a method for
2 managing storage of data in a plurality of storage devices, each storage device having a
3 plurality of blocks for storing data, ~~the method comprising the steps of~~ :

4 generating block layout information in a file system layer of the storage operating
5 system by determining which blocks within the plurality of blocks are allocated for stor-
6 ing data and which are unallocated;

7 transferring the block layout information from the file system layer to a RAID
8 layer of the storage operating system; and

9 responsive to the block layout information, the RAID layer controlling the execu-
10 tion of I/O operations by identifying blocks within the plurality of blocks for use by the
11 I/O operations so as to substantially maximize chain lengths of reads for calculation of
12 parity; ~~and of writes for data write operations;~~

13 thereof while substantially minimizing cost of calculation of parity,

14 wherein ~~the block identification comprises the steps of~~

15 ~~examining all the I/O operations,~~

16 selecting a method for parity calculation which requires the fewest number of read
17 operations to compute substantially minimizes cost of the parity calculation for the
18 I/O operations; and

19 responsive to the block layout information and the parity calculation method se-
20 lection selected, identifying the blocks within the plurality of blocks for use by the I/O
21 operations.

1 2. (Currently Amended) A method for managing storage of data in a plurality of storage
2 devices, each comprising a plurality of storage blocks, ~~the method~~ comprising ~~the steps~~
3 ~~of~~ :

4 generating block layout information; and

5 in response to the block layout information, controlling the execution of an I/O op-
6 eration by identifying storage blocks for use by the I/O operation so as to substantially
7 minimize the number of read operations needed for eost of calculation of error correc-
8 tion parameters across a stripe.

1 3. (Original) The method of claim 2 wherein the calculation of error correction parame-
2 ters comprises the calculation of parity.

1 4. (Original) The method of claim 3 wherein the calculation of parity comprises selecting
2 a parity calculation operation from a group consisting of a subtraction method and a par-
3 ity re-calculation method.

1 5. (Original) The method of claim 2 wherein the identification of storage blocks for use
2 in the I/O operation substantially maximizes the chain length by substantially maximizing
3 the number of blocks having a contiguous physical layout on a storage device.

1 6. (Currently Amended) The method of claim 2, further comprising: ~~the step of~~

2 identifying storage blocks for use in the I/O operation so as to substantially
3 maximize the chain length by substantially maximizing the number of blocks having se-
4 quential VBN's associated with the storage blocks.

1 7. (Currently Amended) The method of claim 2, further comprising: ~~the step of~~
2 identifying storage blocks for use in the I/O operation so as to substantially
3 maximize the chain length by substantially maximizing the locality of the blocks of a
4 storage device.

1 8. (Currently Amended) The method of claim 2 wherein the execution controlling step
2 comprises ~~the steps of~~ :

3 examining blocks to which data is to be written prior to write operations;
4 selecting one of a plurality of parity calculation methodologies including,
5 a first methodology comprising minimizing the number of blocks read, and
6 a second methodology comprising maximizing chain lengths of blocks read for
7 the parity calculation.

1 9. (Currently Amended) The method of claim 8, wherein the execution controlling step
2 further comprises ~~the steps of~~ :

3 implementing the selection responsive to the block layout information; and
4 wherein, if the selection constitutes substantially minimizing the number of read
5 blocks,
6 determining on a stripe-by-stripe basis whether to calculate parity based on a sub-
7 traction method or a recalculation method,

8 performing any appropriate read operations to support the method selected, and
9 calculating parity responsive to the read blocks and the data to be written; and
10 wherein, if the selection constitutes substantially maximizing chain lengths of blocks
11 read,

12 deciding which storage blocks to read to substantially maximize chain length while
13 substantially minimizing the number of storage blocks read to support either a subtraction
14 method or a recalculation method,

15 performing read operations on the blocks to be read, and
16 calculating parity responsive to the read blocks and the data to be written.

1 10. (Original) The method of claim 2, wherein the identification of storage blocks is
2 based at least in part on an available resource.

1 11. (Original) The method of claim 2 further comprising transmitting the block layout
2 information from a file system layer to a RAID layer.

1 12. (Original) The method of claim 2 wherein the generating step further comprises:
2 making a first determination as to whether a storage block is unallocated;
3 making a second determination as to a current implementation of the plurality of
4 storage devices; and
5 generating the block layout information based at least in part on the first and the sec-
6 ond determinations.

1 13. (Original) The method of claim 2, wherein the I/O operation is one of a plurality of
2 I/O operations and one of the plurality of I/O operations is a read operation.

1

2 14. (Original) The method of claim 2, wherein the chain length is a chain length of a read
3 operation for calculation of parity.

1 15. (Original) The method of claim 2, wherein the chain length is a chain length for a
2 write operation for the data.

1 16. (Currently Amended) A method for managing storage of data in a storage system
2 comprising:

3 maintaining a plurality of storage devices each having comprising a plurality of
4 storage blocks; ~~, the method comprising~~

5 writing data to predetermined storage blocks across a plurality of stripes and to
6 predetermined storage blocks within each storage device so as to substantially maximize
7 chain length of storage blocks within each storage device while minimizing a number of
8 read operations for the ~~substantially minimizing cost of~~ calculation of error correction
9 parameters across each stripe of the plurality of stripes.

17. - 38. (Cancelled)

1 39. (Currently Amended) A storage system comprising:

2 a plurality of storage devices each having comprising: a plurality of storage
3 blocks; and

4 a storage manager in communication with the plurality of storage devices, the stor-
5 age manager writing data to predetermined storage blocks across a plurality of stripes and
6 to predetermined storage blocks within each storage device so as to substantially maxi-
7 mize chain length of storage blocks within a storage device while substantially minimiz-
8 ing a number of read operations required for calculation the calculation of error correc-
9 tion parameters across each stripe of the plurality of stripes.

1 40. (Currently Amended) A system for managing the storage of data, the system
2 comprising:
3 a plurality of storage devices each having a plurality of storage blocks;
4 a storage device manager in communication with the plurality of storage blocks;
5 a block layout information generator in communication with the storage device man-
6 ager and the plurality of storage blocks; and
7 an error correction parameter calculator in communication with the plurality of stor-
8 age blocks and the storage device manager,
9 wherein the storage device manager, in response to the block layout information
10 from the block layout information generator, controls the execution of an I/O operation
11 by identifying storage blocks for use by the I/O operation so as to substantially maximize
12 chain length within the storage device while substantially minimizing the number of read
13 operations required for calculation by the error correction parameter calculator of error
14 correction parameters across a stripe.

1 41. (Currently Amended) A method for managing storage of data in storage blocks,
2 the method comprising ~~the steps of~~ :
3 generating block layout information;
4 dynamically determining a first number of error correction calculations;

5 dynamically determining a second number corresponding to a chain length; and
6 in response to the block layout information, controlling the execution of an I/O op-
7 eration by identifying storage blocks for use by the I/O operation so as to have a chain
8 length of the second number within a storage device while performing the first number of
9 calculations of error correction parameters across a stripe.

42. - 44. (Cancelled)

Please add new claims 45, *et seq.*, as follows:

1 45. (New) A method for managing storage of data by a server, comprising:

2

3 receiving a request to write data to a plurality of storage devices;

4

5 generating block layout information to determine which blocks within a plurality
6 of blocks located in the plurality of storage devices are allocated for storing data and
7 which are unallocated;

8

9 identifying blocks within the plurality of blocks for use by a set of I/O operations
10 to store the data;

11

12 determining the number of read operations needed to compute parity for the data
13 by computing parity using a subtraction method of computing parity;

14

15 determining the number of read operations needed to compute parity for the data
16 by computing parity using a recalculation method of computing parity;

17

18 choosing either the subtraction method of computing parity or the recalculation
19 method of computing parity by determining which of these two methods requires the
20 fewer number of read operations, and choosing the method requiring the fewer number of
21 read operations; and

22

23 writing the data to identified blocks, and computing parity for the data using the
24 chosen method of computing parity.

1 46. (New) The method of claim 45, further comprising:

2 choosing to either firstly maximize chain lengths of read operations for calcula-
3 tion of parity or secondly to place the data with a high degree of locality in the plurality
4 of storage devices, by choosing the method which requires the fewest number of read op-
5 erations in computing parity for the data.

1 47. (New) A method for managing storage of data by a server, comprising:
2 receiving a request to write data to a plurality of storage devices;
3 generating block layout information to determine which blocks within a plurality
4 of blocks located in the plurality of storage devices are allocated for storing data and
5 which are unallocated;
6 identifying blocks within the plurality of blocks for use by a set of I/O operations
7 to store the data; and
8 selecting whether to substantially minimize the number of read blocks or to sub-
9 stantially maximize chain lengths of read blocks, and implementing the selection respon-
10 sive to the block layout information, and responsive to the method requiring the fewer
11 number of read operations.

1 48. (New) The method of claim 47, further comprising:
2 in response to selecting to substantially minimize the number of read blocks, de-
3 termining whether to calculate parity based on the subtraction method or the recalculation
4 method by determining which method requires the fewer number of read operations, and
5 selecting the method which requires the fewer number of read operations,
6 performing the write operation and calculating the parity using the parity calcula-
7 tion method requiring the fewer number of read operations.

1 49. (New) The method of claim 48, further comprising:
2 in response to selecting to substantially maximize chain lengths of read blocks,
3 deciding which storage blocks to read to substantially maximize chain length while
4 minimizing the number of storage blocks read to support either the subtraction method or
5 the recalculation method of parity calculation; and

6 performing the write operation and calculating the parity using the parity calcula-
7 tion method requiring the fewer number of read operations.

1 50. (New) A method for managing storage of data by a server, comprising:
2 receiving a request to write data to a plurality of storage devices;
3 generating block layout information to determine which blocks within a plurality
4 of blocks located in the plurality of storage devices are allocated for storing data and
5 which are unallocated;
6 identifying blocks within the plurality of blocks for use by a set of I/O operations
7 to store the data;
8 testing to either maximize chain lengths of read operations for calculation of par-
9 ity, or to place the data with a high degree of locality in the plurality of storage devices,
10 the testing having the steps,
11 determining, for both maximizing chain length and placing the data with a high
12 degree of locality, the number of read operations needed to compute parity for the data,
13 by computing parity using both the subtraction method of computing parity and the recal-
14 culation method of computing parity;
15 firstly choosing to either maximize chain lengths of read operations for
16 calculation of parity or to place the data with a high degree of locality in the plu-
17 rality of storage devices, and after this first choice, secondly choosing either the
18 subtraction method of computing parity or the recalculation method of computing
19 parity by determining which of these methods requires the fewest number of read
20 operations,
21 choosing the method requiring the fewest number of read operations of
22 computing parity of the data; and
23 writing the data to identified blocks, and computing parity for the data using the
24 chosen method of computing parity.

1 51. (New) A computer readable media, comprising:

2 said computer readable media containing instructions for execution on a processor
3 for a method of managing storage of data in a plurality of storage devices, each storage
4 device having a plurality of blocks for storing data, the method having,
5 generating block layout information; and
6 in response to the block layout information, controlling the execution of an I/O op-
7 eration by identifying storage blocks for use by the I/O operation so as to substantially
8 minimize the number of read operations needed for calculation of error correction pa-
9 rameters across a stripe.